David B. Shaffer: RFI: Effects on Bandwidth Synthesis, IVS 2000 General Meeting Proceedings, p.402–406
http://ivscc.gsfc.nasa.gov/publications/gm2000/shaffer

RFI: Effects on Bandwidth Synthesis

David B. Shaffer

RadioMetrics/NVI, Inc./NASA Goddard Space Flight Center
e-mail: dbs@gemini.gsfc.nasa.gov

Abstract

I describe briefly how bandwidth synthesis works for VLBI, and then show how specific levels of RFI, expressed as a percentage increase above the nominal SEFD, affect the measured group delay.

1. Introduction

This note is a result of discussions held during an RFI workshop which took place at the Wettzell site on February 24, 2000 as part of the IVS meeting.

In the previous contribution, Brian Corey described radio frequency interference (RFI) measurements and how to determine the RFI power that may be incident on an antenna used for VLBI measurements.

The intent here is to describe quantitatively how RFI affects VLBI geodesy measurements and to quantify the level of RFI which is harmful.

2. Group Delay and Bandwidth Synthesis (BWS)

The most important quantity determined for a VLBI geodesy scan is the "group delay". This delay is the relative difference in the time of arrival of the random noise "signal" from a radio source at two antennas. As the Earth turns, the delay changes continuously, in a manner which depends on the relative location of the antennas and the position of the radio source. (Although many antennas participate in most VLBI experiments, in Mk III/Mk IV data analysis, the data from each pair of antennas is treated independently.)

Our ability to determine an accurate value for the group delay (at S- or X-band) depends on how much bandwidth is analyzed. Consider a wide bandwidth "signal" which is basically random noise. (This is what comes from the quasar radio sources.) The signal results from the addition of all the random electromagnetic fields within the overall bandwidth. This superposition of signals changes significantly on a time scale which depends on the maximum frequency difference (that is, the bandwidth). When we compare the two versions of this signal sampled (recorded) at different sites (this is what the correlator does!), we will find a non-zero correlation only if the two versions are closely aligned in time.

If B is the bandwidth that we process, the estimate of the group delay at which the signals are best aligned has an accuracy of about $1/B$. If the signal-to-noise ratio (SNR) is high, we can tell more accurately when the signals are aligned. In fact, the actual error (σ) in determining the best alignment delay is $\sigma = 1/(2\pi \cdot B_{rms} \cdot SNR)$. B_{rms} is the root-mean-square spanned bandwidth, which is about 40% of the total frequency span for the frequency sequences that we use.

How accurately must we measure the group delay?, *i.e.* what are the requirements on σ? We are trying to measure geodetic properties with a precision of better than a centimeter. Light travels one centimeter in 33 trillionths of a second: $33 \cdot 10^{-12}$ seconds. One trillionth of a second is called a picosecond. Thus, we need to make group delay measurements with an accuracy of tens of picoseconds in order to achieve our desired geodetic goals. If we can achieve a modest SNR of 20, then we must analyze signals with bandwidths of hundreds of megahertz.

If we could record a complete bandwidth of several hundred MHz, our measurement of group delay would be relatively simple. We would just adjust the relative timing of the tapes until the correlator produced the maximum cross-correlation signal. However, the large bandwidth would require a very high data recording rate: twice the maximum bandwidth, or a sample rate approaching a gigabit per second. Although this rate should be achievable with the Mk IV, it is much higher than previous VLBI systems could record. Hence, a better way to sample a broad bandwidth was needed. Haystack's Alan Rogers showed in a 1970 paper that you don't need to record the entire bandwidth ("Very Long Baseline Interferometry with Large Effective Bandwidth for Phase Delay Measurements," *Radio Science*, **5**, 1239-1247). He demonstrated how to achieve nearly the same result by recording several narrow frequency channels spread out across the desired band. This technique, called "Bandwidth Synthesis" (BWS), is how we now record geodetic VLBI data, and is the reason for all the video converters in the data acquisition systems.

We must be careful how we combine the data from the BWS channels. In the case of a single very broad bandwidth, when the tapes are correlated at exactly the right delay, all parts of the processed bandwidth are in phase, with zero relative offset. If the estimated delay is not quite right, there will be a linear phase shift across the bandwidth. This shift arises because the real signal travel time between the two antennas does not match the delay picked by the correlator. In effect, the signal went a little farther (or a little less far) than the correlator estimated. At the higher frequencies in the band, the signal wavelength is shorter. Thus, at the higher frequencies, the extra distance that the signal went corresponds to more cycles (more phase) of the wavelength. If we can measure the shift of phase with frequency (the "phase slope") across the bandwidth, we can determine the difference between the delay used by the correlator (the "model delay") and the true delay. By adding the delay corresponding to the phase slope to the model delay, we get the actual delay. A set of such delays, determined for many scans, is then used to determine the overall geometry of the observations.

The BWS process takes advantage of this phase slope to measure the group delay correction to the correlator model: we process all the frequency channels, look at how the phase changes from channel to channel, and use this change to determine the delay.

This means we have to know how to align the various frequency channels. That is the purpose of the phase calibration system. The phase cal signals are injected in phase at the feed of the receiver. Any channel-to-channel variation of the phase cal phases when extracted from the individual frequency channels is due to differing path lengths through the VLBI equipment and differing video converter LO phases. The measured calibration phases are applied to the measured signal phases to take out the channel offsets. If the phase calibration fails, we can also observe a strong source to determine directly the phase offsets and apply the relative offsets to subsequent data: so-called "manual phase cal." This latter technique, however, does not let us keep track of time-dependent instrumental delay changes, so we really do prefer that the phase cal work!

Figure 1 shows a set of (fake!) BWS data, for our normal X-band observation mode, where eight frequency channels are used. The spacing of the channels is carefully chosen to maximize the effective spanned bandwidth and minimize confusing effects from combining the channels. At X-band, the channels that we use are often arranged in the sequence 0-1-4-10-21-29-34-36. These

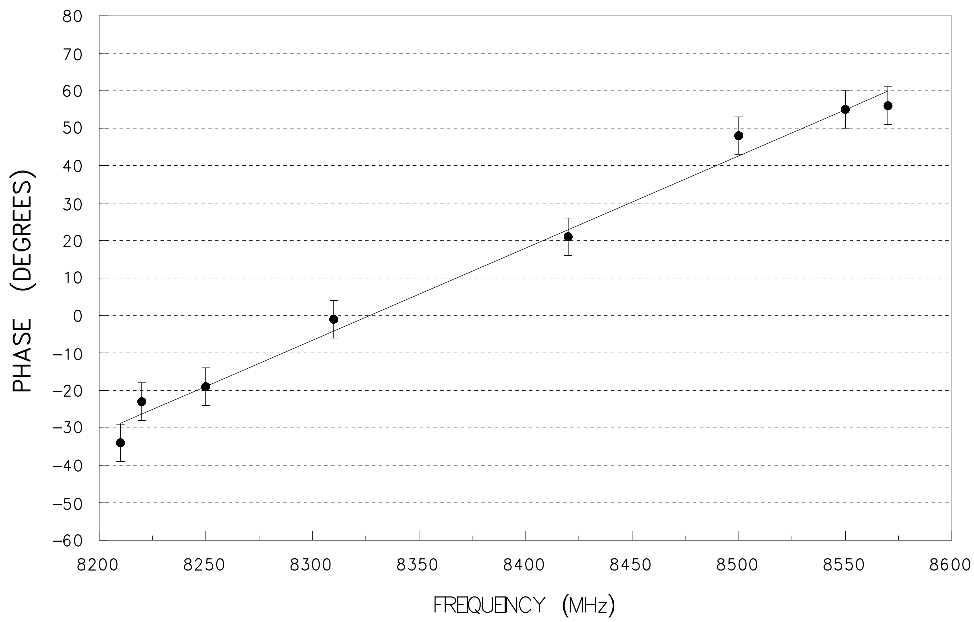

Figure 1. A (hypothetical) plot of phase versus frequency for an X-band experiment using the CDP standard narrowband sequence.

values denote the relative spacing between the channels. This sequence is usually multiplied by either 10 MHz or 20 MHz.

Because our equipment is not perfect and the phase calibration is not perfect, either, the amplitude of the correlated signal in the various channels is generally not the same, and there are residual phase offsets in each channel, too, as shown in Figure 1. Typical phase offsets are several degrees, resulting from such problems as phase offsets in the feeds, reflections in the phase cal system, and spurious phase cal signals. The amplitude and phase imperfections cause a group delay offset. As long as the phase offsets and the relative signal strength in each channel do not change, the group delay offset remains constant, and does not affect geodetic determinations. The offset is included in the clock offset term determined for each baseline.

3. RFI Effects

If the channel offsets change in a random sense during an experiment, there *will be* additional random group delay errors. If there are systematic changes, the delays will also be affected systematically. This is how RFI can cause serious problems.

Consider Figure 1. Determining the group delay is equivalent to fitting a line through the (frequency, phase) values for all the points. The slope of that line gives the group delay. Obviously, if the phase changes, a different slope (different delay) results. Generally, RFI does not cause phases to change. (However, some kinds of RFI—coherent signals at the phase cal frequencies—can cause the phase calibration phase to change, which will cause corresponding phase errors when the calibration is applied.) There is an error, however, on the measurement of the phase in each BWS channel. This error depends directly on the SNR in the channel. The size of this phase error *does* affect the fitting of the line through the phases: if the phase error in some channel is larger (because the SNR is lower), that channel has a reduced effect on the fit. In the worst case, there may be so much RFI that a particular channel must be deleted from the fit. In this case, the effect of RFI is pretty obvious.

We can quantify the effect of RFI by noting how it affects the SNR in a frequency channel, and then how that affects the group delay determination. The channels for which RFI has the most effect are the end channels of the frequency sequences. These channels have the most leverage on the fit for the group delay. Since it is the overall SNR for the baseline (both antennas used together) that counts, we must relate the RFI level at one antenna to its effect on the interferometer.

The signal-to-noise ratio (SNR) of a VLBI observation depends on three parameters:

- the correlated flux density of the radio source: S_c
- the System Equivalent Flux Densities (SEFD) of the antennas
- the total number of bits correlated: N

Note that N depends on the scan length and the sample rate.

The exact expression for SNR is $SNR = S_c\sqrt{N}/\sqrt{SEFD_1 \cdot SEFD_2}$.

I now evaluate what happens to the group delay when a scan is affected by RFI at only one antenna. I assume that the correlated flux density, the SEFD at the second antenna, and the number of bits correlated do not change. Basically, the RFI increases the SEFD at the first antenna, which in turn reduces the overall SNR for the scan. The SEFD depends on the size of the antenna and the system noise power. (This power is usually expressed in terms of an equivalent noise temperature. Note that "temperature" and "power" are essentially equivalent, related by Boltzmann's equation: $P = kTB$.) When there is no RFI, the system noise temperature is determined by the internal noise power generated in the receiver amplifier, as well as the addition of some external radiation from the atmosphere and the ground. (Strictly speaking, the radio source itself also contributes to the system noise, but since most of the sources we observe add much less than one percent to the overall noise power, we usually neglect their noise contribution.) When there is RFI, the RFI power adds directly to the system noise power to increase the SEFD.

Thus, we can relate RFI to SEFD and then SNR by comparing the system power level with and without the RFI. If the RFI raises the power in a video convertor (as read by the Field System, using the `TPI` command, for example) by 10%, the effective SEFD in that channel will also increase by 10%, or a factor of 1.1, and the baseline SNR will be reduced by a factor of $1/\sqrt{1 \cdot 1.1} = 0.95$. If the RFI power is 100% of SEFD, it will double the power level, and the SEFD will also be a factor of 2 higher. Then the baseline SNR will be a factor of $1/\sqrt{1 \cdot 2} = 0.707$ lower.

Tables 1 and 2 show the effect RFI can have on group delay measurements when a single channel has a phase offset of 5 degrees, an offset which is not at all unusual. The RFI levels at one antenna are expressed as a fractional addition to the nominal system power level *in the affected channel only.* The "Relative Baseline SNR" shows the reduction in SNR, but only for the affected channel. I have applied the offset to the end channel of the frequency sequence to show the maximum effect. These tables show the group delay offset compared to the case of no phase error in the affected channel. The effect of RFI is to *reduce* the offset. This reduction, though, is not good—the original offset, without RFI, would just be absorbed into a clock offset. With RFI, what we consider in our analysis to be a clock offset is not constant!

The effect of RFI seems much worse at S-band. This is because the total frequency span is not as wide. In this case, a given phase offset has a larger effect on the slope of the phase line. Fortunately, the size of S-band delay errors is reduced by a factor of about 13 in our data analysis when the ionosphere correction is applied. (This correction combines the delays measured at S-and X-band. We observe at two frequencies just so we can make this correction.)

Table 1. Effects of RFI at X-band

Single Antenna RFI Level	Relative Baseline SNR	Group Delay Offset
No RFI	1.000	16.9 picosec
10% RFI	0.953	15.9
20% RFI	0.913	15.0
30% RFI	0.977	14.2
40% RFI	0.845	13.4
50% RFI	0.816	12.8
100% RFI	0.707	10.3

Frequency Sequence is 0-1-4-10-21-29-34-36, multiplied by 10 MHz
RFI and a 5° phase offset occur in channel 8 only (frequency spacing 360 MHz)

Table 2. Effects of RFI at S-band

Single Antenna RFI Level	Relative Baseline SNR	Group Delay Offset
No RFI	1.000	96.9 picosec
10% RFI	0.953	92.1
20% RFI	0.913	87.9
30% RFI	0.977	83.9
40% RFI	0.845	80.3
50% RFI	0.816	77.1
100% RFI	0.707	64.0

Frequency Sequence is 0-1-4-10-15-17, multiplied by 5 MHz
RFI and a 5° phase offset occur in channel 6 only (frequency spacing 85 MHz)

At X-band, moderate RFI levels (those which increase the system noise power in one channel by less than 50%) can easily cause delay errors of several picoseconds, or more than 1 mm of geometric error. Larger RFI levels, of course, result in even larger errors, compared to the situation when there is no RFI. Notice that varying RFI (the usual case - RFI is seldom constant!) really causes a varying delay bias, rather than delay noise which could be either positive or negative. Thus, RFI will tend to "pull" the geodetic results, rather than just making the results noisier. For instance, if RFI is worse in a particular direction (a typical situation), the delays measured when the antenna is pointed in that direction will be affected systematically, leading to a biased position for the antenna. This biasing is probably the most serious reason why RFI is undesirable.

Based on simulations such as those used to generate the Tables, we have chosen an RFI level which causes a 10% increase in system noise power in a video converter bandwidth to be the level at which we begin to worry about degradation of VLBI observations. This is a quantifiable standard which we can present to other agencies as well as entities which are potential generators of RFI.